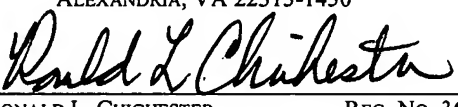


CERTIFICATE OF MAILING via EXPRESS MAIL 37 C.F.R. 1.10	
PURSUANT TO 37 C.F.R. §1.10, I HEREBY CERTIFY THAT I HAVE A REASONABLE BASIS FOR BELIEF THAT THIS CORRESPONDENCE IS BEING DEPOSITED WITH THE UNITED STATES POSTAL SERVICE EXPRESS MAIL POST OFFICE TO ADDRESSEE ON THE DATE INDICATED BELOW, ADDRESSED TO:	
MAIL STOP PATENT APPLICATION HONORABLE COMMISSIONER FOR PATENTS P.O. Box 1450 ALEXANDRIA, VA 22313-1450	
	
RONALD L. CHICHESTER	REG. No. 36,765
DATE OF MAILING:	JANUARY 30, 2004
EXPRESS MAIL LABEL:	EV339225060US

APPLICATION FOR LETTERS PATENT

FOR

AUTOMATIC MEDIA REPAIR AFTER READ FAILURE DUE TO MEDIA ERROR

INVENTORS:	Ahmad A.J. Ali
ASSIGNEE:	Dell Products L.P.
ATTORNEY:	Ronald L. Chichester of Baker Botts L.L.P.
ATTORNEY DOCKET NO.:	016295.1522
CLIENT REFERENCE:	DC-05817/AEP

AUTOMATIC MEDIA REPAIR AFTER READ FAILURE DUE TO MEDIA ERROR**Background**Field of the Disclosure

[0001] The present disclosure relates in general to the field of data storage systems and, more particularly, to a system and method for repairing, in an automated fashion, the media of the storage system after an error is encountered in the media.

Background of the Related Art

[0002] As the value and the use of information continue to increase, individuals and businesses seek additional ways to process and store information. One option available to users is information handling systems. An information handling system generally processes, compiles, stores and/or communicates information or data for business, personal or other purposes, thereby allowing users to take advantage of the value of the information. Because technology and information handling needs and requirements vary between different users or applications, information handling systems may also vary regarding what information is handled, how the information is handled, how much information is processed, stored, or communicated, and how quickly and efficiently the information may be processed, stored, or communicated. The variations in information handling systems allow for information handling systems to be general or configured for a specific user or specific use such as financial transaction processing, airline reservations, enterprise data storage, or global communications. In addition, information handling systems may include a variety of hardware and software components that may be configured to process, store, and communicate information and may include one or more computer systems, data storage systems, and networking systems, *e.g.*, computer, personal computer workstation, portable computer, computer server, print server, network router, network

hub, network switch, storage area network disk array, redundant array of independent disks (“RAID”) system and telecommunications switch.

[0003] Computer systems often include hard media, such as IDE and/or SCSI devices. Hard media errors during read operations on SCSI drives under RAID controllers are gracefully handled for redundant RAID configurations (such as in RAID levels 1, 5, or 10) but not on non-redundant configurations (such as RAID level 0, or degraded levels 1, 5, or 10) where there is no recovery mechanism. The host level software application may experience a read failure when a media error is encountered because the data associated with software application is stored at the location of the media error and is thus inaccessible and/or corrupted.

[0004] One problem scenario is when a user attempts to restore data from a backup. Part of the restored data may again be written to the same (bad) sector that caused the read error originally. SCSI drives do not track sectors that have caused read errors previously, and new write commands to the bad sector may be completed without any verification and thus reported as being completed successfully. Subsequent read commands from that bad sector may result in an unrecoverable error due to lack of data availability or corruption.

[0005] A second problem scenario is when a user performs a “verify” operation on the SCSI disk. In that case, the verify operation would detect the bad sector on the disk and reassign a good sector (from the spare sectors) in place of the bad sector. The problem with this operation is that unknown “data” (in the form of “1’s and 0’s”) exists on the newly assigned good sector. The software application that was using the data on the bad sector is unaware of the reassignment by the verify operation, and hence does not know that a block of data (from the bad sector) is now of unknown status or validity. Indeed, the software application could issue a read request for the data in the reassigned sector and inadvertently read the unknown data that was present in

the new sector when it was reassigned during the verify operation. The software application would then be working on unknown, and potentially corrupted data, which may result in a crash of the software application, or produce inaccurate results. A user may restore the damaged file after the repair, but the verify operation may have reassigned/repared other bad sectors that were discovered during the verify operation and the files residing on those sectors would (presumably) be corrupted. Moreover, the files in question may have already been corrupted (due to a bad sector) but went unnoticed because those sectors had not undergone a read operation.

[0006] In the past, recovery from media errors on SCSI drives required a complete restore operation from backup (assuming that a backup existed). A complete recovery was warranted because it was hard to determine which files were corrupted and/or damaged due to bad sectors that were uncovered during the verify operation. There is, therefore, a need in the art for a system and/or method for avoiding bad sectors on a storage media while maintaining operation of that media.

Summary of the Invention

[0007] In accordance with the present disclosure, a system and method are provided that performs automatic media repair so that, after a media error is encountered, subsequent write operations are completed on a known good sector while read operations to the repaired sector are induced to fail so that the user never receives corrupted or undetermined data. Another advantage of the present disclosure is that recovery of lost data due to the media error is achievable by restoring only the damaged file, rather than restoring the complete media volume. Consequently, recovery is quicker and only affects one of the working processes, rather than the system as a whole.

[0008] This disclosure provides a method for media repair of a storage device, wherein the storage device performs a read operation on the storage device, detects a read error, locks a logical block address on the storage device, performing a reassign operation on the storage device, performs a write (signature and date) operation on the storage device, and unlocks the logical block address after the write operation. Alternatively, if an error is not detected from the read operation, the method may perform detect a signature (such as an ECC signature), and perform a write operation on the storage device. If no signature is found, the method can lock a logical block address on the storage device and perform a write operation on the storage device to place the signature, and unlock the logical block address. The storage device can be in a non-RAID or non-redundant RAID configuration. Moreover, to facilitate the method disclosed herein, the read operation is a READ LONG operation, the write operation is a WRITE LONG operation which may produce invalid ECC data. The storage device in question can be a SCSI device, and IDE device, an ATA device, or similar.

[0009] Other technical advantages should be apparent to one of ordinary skill in the art in view of the specification, claims, and drawings.

Brief Description of the Drawings

[0010] A more complete understanding of the present disclosure and advantages thereof may be acquired by referring to the following description taken in conjunction with the accompanying drawings, in which like reference numbers indicate like features, and wherein:

[0011] Figure 1 depicts a component diagram of a storage area network including one embodiment of a resource management engine that incorporates the teachings of the present disclosure;

[0012] Figure 2 is a block diagram illustrating a mass storage device having a sector and a bad section of that sector;

[0013] Figure 3 is a flowchart of an embodiment of the present disclosure.

[0014] The present disclosure may be susceptible to various modifications and alternative forms. Specific exemplary embodiments thereof are shown by way of example in the drawing and are described herein in detail. It should be understood, however, that the description set forth herein of specific embodiments is not intended to limit the present disclosure to the particular forms disclosed. Rather, all modifications, alternatives, and equivalents falling within the spirit and scope of the invention as defined by the appended claims are intended to be covered.

Detailed Description of the Preferred Embodiments

[0015] The present disclosure provides a system and method for a RAID controller or a non-RAID controller that performs automatic media repair so that, after a media error is encountered, subsequent write operations are completed on a known good sector while read operations to the repaired sector are induced to fail so that the user never receives corrupted or undetermined data. Another advantage of the present disclosure is that recovery of the error is only by restoring the damaged file, rather than restoring the complete media volume. Consequently, recovery is quicker and only affects one of the working processes, rather than the system as a whole. The method disclosed herein is particularly useful for non-RAID and non-redundant RAID configurations.

[0016] In one embodiment, a method is employed that uses SCSI REASSIGN, WRITE LONG, and READ LONG commands. The REASSIGN command allows the disk to remap the bad sector into a reserved sector. The WRITE LONG command allows the manipulation of error

checking and correction (“ECC”) data for the reassigned sector so that the data associated with the bad sector appears to be corrupted to a READ command, but would still allow a WRITE command to complete with proper ECC data for that bad sector.

[0017] In another embodiment, the controller is able to differentiate between repaired sectors (having one or more soft read errors) and unrepaired bad sectors (those with one or more hard error) by using READ LONG commands based on a signature that is written on the repaired sectors using the WRITE LONG command.

[0018] In another embodiment, a counter and date can be stored along with the signature on each repaired sector in order to avoid multiple event logging and/or user notification for a single sector. Moreover, the technique can be used to track the age of the repaired (but not corrected) sector.

[0019] While the embodiments above utilized commands according to the SCSI standard, other disk drives, such as integrated drive electronics (“IDE”) devices and/or advanced technology attachment (“ATA”) devices and/or regular IDE drives may benefit from the method disclosed herein. In order to work according to the method disclosed herein, however, the alternate drive type would have to support something like the WRITE LONG command, although the device in question need not support the exact same WRITE LONG as the SCSI specification. For example, something like the SoftCorruptBlock command could be used in order to implement the method disclosed herein.

[0020] Elements of the present disclosure can be implemented on a computer system, as illustrated in Figure 1. Referring to Figure 1, depicted is an information handling system, generally referenced by the numeral 100, having electronic components mounted on at least one printed circuit board (“PCB”) (not shown) and communicating data and control signals

therebetween over signal buses. In one embodiment, the information handling system is a computer system. The information handling system comprises processors 110 and associated voltage regulator modules (“VRMs”) 112 configured as processor nodes 108. There may be one or more processor nodes 108, one or more processors 110, and one or more VRMs 112, illustrated in Figure 1 as nodes 108a and 108b, processors 110a and 110b and VRMs 112a and 112b, respectively. A north bridge 140, which may also be referred to as a “memory controller hub” or a “memory controller,” is coupled to a main system memory 150. The north bridge 140 is coupled to the processors 110 via the host bus 120. The north bridge 140 is generally considered an application specific chip set that provides connectivity to various buses, and integrates other system functions such as memory interface. For example, an INTEL® 820E and/or INTEL® 815E chip set, available from the Intel Corporation of Santa Clara, California, provides at least a portion of the north bridge 140. The chip set may also be packaged as an application specific integrated circuit (“ASIC”). The north bridge 140 typically includes functionality to couple the main system memory 150 to other devices within the information handling system 100. Thus, memory controller functions, such as main memory control functions, typically reside in the north bridge 140. In addition, the north bridge 140 provides bus control to handle transfers between the host bus 120 and a second bus(es), *e.g.*, PCI bus 170 and AGP bus 171, the AGP bus 171 being coupled to the AGP video 172 and/or the video display 174. The second bus may also comprise other industry standard buses or proprietary buses, *e.g.*, ISA, SCSI, USB buses 168 through a south bridge (bus interface) 162. These secondary buses 168 may have their own interfaces and controllers, *e.g.*, RAID Array storage system 160 and input/output interface(s) 164. Finally, a BIOS 180 is operative with the information handling system 100 as illustrated in Figure 1. The information handling system 100 can be combined

with other like systems to form larger systems. Moreover, the information handling system 100, can be combined with other elements, such as networking elements, to form even larger and more complex information handling systems.

[0021] Figure 2 illustrates a mass storage device 200, such as a SCSI device, having a storage disk 202 that has at least one sector 204. The storage device 200 can be a standalone device, or be part of the RAID array 160 (see Figure 1). In this illustration, a file 206 is stored within the sector 206. A bad portion 208 of the sector 204 can arise from any number of factors as is commonly experienced in the art. As the bad portion 208 arose within the space allocated for the file 206, one or more bytes of the file 206 may be corrupted or indeterminate. Hence, some error correction mechanism is needed to ensure integrity of the file, preferably without removing the disk 202 from operation.

[0022] Figure 3 illustrates the method of the present disclosure. The media READ error recovery method begins generally at step 302. First, in step 304, a READ LONG command is issued to the disk 200. A check is then made in step 306 to determine if an error was encountered by performing step 302. If an error was encountered (*i.e.*, the result of step 306 is positive) then the logical block address (“LBA”) for the device 200 is locked. Thereafter, in step 310, a REASSIGN command is issued to the device 200. Then, in step 312, a WRITE LONG command is issued with *invalid* ECC data, preferably with a signature, counter, and date information. Once the WRITE LONG operation has been completed, the logical block address can be unlocked in step 314, and the method ends generally at step 340.

[0023] An alternate scenario occurs when a READ LONG operation does not produce an error (*i.e.*, the result of step 306 is negative). In that case, a check is made in step 320 to determine if a signature is found as a result of the READ LONG command. If a signature was

found (*i.e.*, the result of step 320 is positive), then a WRITE LONG command is executed with invalid ECC data, and a counter is increased, after which the method ends generally at step 340.

[0024] Another alternate scenario occurs when a signature was not found (*i.e.*, the result of step 320 is negative). In that case, the LBA is locked in step 330. Next, in step 332, the WRITE LONG command is executed with invalid ECC data (including signature, counter, and date information). Once the WRITE LONG command has been completed, the LBA is unlocked in step 334 and the method ends generally at step 340.

[0025] The invention, therefore, is well adapted to carry out the objects and to attain the ends and advantages mentioned, as well as others inherent therein. While the invention has been depicted, described, and is defined by reference to exemplary embodiments of the invention, such references do not imply a limitation on the invention, and no such limitation is to be inferred. The invention is capable of considerable modification, alteration, and equivalents in form and function, as will occur to those ordinarily skilled in the pertinent arts and having the benefit of this disclosure. The depicted and described embodiments of the invention are exemplary only, and are not exhaustive of the scope of the invention. Consequently, the invention is intended to be limited only by the spirit and scope of the appended claims, giving full cognizance to equivalents in all respects.